

Cosmological constraints from the Subaru Hyper Suprime-Cam year 1 shear catalogue lensing convergence probability distribution function

Leander Thiele^{1,2,*}, Gabriela A. Marques^{3,4,5}, Jia Liu,² and Masato Shirasaki^{6,7}

¹*Department of Physics, Princeton University, Princeton, New Jersey 08544, USA*

²*Center for Data-Driven Discovery, Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan*

³*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

⁴*Kavli Institute for Cosmological Physics, University of Chicago, Chicago, Illinois 60637, USA*

⁵*Florida State University, 77 Chieftan Way, Tallahassee, Florida 32306, USA*

⁶*National Astronomical Observatory of Japan (NAOJ), National Institutes of Natural Sciences, Osawa, Mitaka, Tokyo 181-8588, Japan*

⁷*The Institute of Statistical Mathematics, Tachikawa, Tokyo 190-8562, Japan*



(Received 18 April 2023; accepted 20 November 2023; published 13 December 2023)

We utilize the probability distribution function (PDF) of normalized convergence maps reconstructed from the Subaru Hyper Suprime-Cam (HSC) year 1 shear catalogue, in combination with the power spectrum, to measure the matter clustering amplitude $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$. The large-scale structure's statistical properties are incompletely described by the traditional two-point statistics, motivating our investigation of the PDF—a complementary higher-order statistic. By defining the PDF over the standard-deviation-normalized convergence map, we are able to isolate the non-Gaussian information. We use tailored simulations to compress the data vector and construct a likelihood approximation. We mitigate the impact of survey and astrophysical systematics with cuts on smoothing scales, redshift bins, and data vectors. We find $S_8 = 0.860_{-0.109}^{+0.066}$ from the PDF alone and $S_8 = 0.798_{-0.042}^{+0.029}$ from the combination of the PDF and power spectrum (68% confidential level (CL)). The PDF improves the power-spectrum-only constraint by about 10%.

DOI: [10.1103/PhysRevD.108.123526](https://doi.org/10.1103/PhysRevD.108.123526)

I. INTRODUCTION

Weak lensing surveys are rapidly catching up with the precision afforded by cosmic microwave background (CMB) experiments [e.g., [1–3]]. Weak lensing directly probes the clustering of matter through statistical measurements of the distortion of galaxy shapes. In recent years, weak lensing surveys have produced increasingly precise measurements of the matter clustering amplitude S_8 [e.g., [4–8]], showing hints of a discrepancy between galaxy lensing and CMB determinations, dubbed the S_8 -tension. Typically, these studies adopt two-point statistics—the two-point correlation function or the power spectrum.

While two-point statistics are sufficient to describe Gaussian fields such as the CMB, the large-scale structure that sources the weak lensing signal has undergone nonlinear growth and is thus far from Gaussian. Therefore, weak lensing non-Gaussian statistics have

been proposed to extract complementary information.¹ Equally importantly, non-Gaussian statistics are affected by systematics differently than two-point statistics. This is especially relevant for the S_8 -tension, since whether its origin stems from new physics or systematics is currently under hot debate. Non-Gaussian statistics will be particularly beneficial for upcoming surveys such as *Rubin* LSST [54], *Euclid* [55], and *Roman* [56]. These surveys will provide a high source density and hence probe deeper into the nonlinear regime.

In this work, we focus on the probability distribution function [PDF, [57–60]] of the lensing convergence map²

¹Popular non-Gaussian statistics include minima, peaks [9–17], voids [18], Minkowski functionals [19–25], Betti numbers [26,27], persistent homology [28,29], wavelets [30], scattering transform [31,32], moments [21,22,33–37], higher-order correlation functions [38–43], density-split statistics [44], and convolutional neural networks [45–47]. A comparison of some of these in a forecast setting for *Euclid* has been carried out by [48], finding about similar performance for many of these statistics, but it is limited to the Fisher approximation. An even more ambitious attempt uses forward-model, field-level methods [e.g., [49–53]].

²Also called the mass map.

*lthiele@princeton.edu

from the Subaru Hyper Suprime-Cam (HSC) year 1 (Y1) data release [61,62]. HSC is currently the deepest large-scale lensing survey and is thus considered a precursor of Stage-IV surveys. The PDF collects the amplitudes—but not shapes—of correlation functions of all orders. Thus, it is a highly non-Gaussian statistics and contains distinct information “far” from the power spectrum. Forecasts have shown promise in tightening the constraints on not only S_8 but also the neutrino mass sum and the dark energy equation of state [63–68]. Formally, the PDF and moments are closely related. Moments can suffer from instabilities due to a few extreme pixels. This implies that, in practice, with a limited range of convergence values included in the PDF’s bins, the PDF can still be complementary to moments.

Our work is the first to obtain cosmological constraints from the lensing PDF with observations. While there exist analytic models of the PDF based on large-deviation theory [69–71] and the halo model [66,72,73], they are not yet at the level of accuracy and flexibility required to incorporate the complex survey configurations and systematics. Therefore, we use a large set of cosmological simulations tailored to HSC Y1 data to model the lensing PDF and its likelihood.

II. METHODS

In this section, we provide a brief overview of the HSC Y1 data, the PDF and power spectrum measurements, the simulations suite, the likelihood, our blinding procedure, and systematics. More detail may be found in our companion paper on counts of lensing peaks and minima [74].

A. HSC Y1 data

We use lensing convergence maps estimated from the HSC Y1 galaxy shapes catalogue [75]. After applying masks, the shear maps span 136.9 deg^2 in six spatially disjoint fields. We adopt galaxy redshifts determined using the MLZ code [76]. After applying redshift cuts of $0.3 < z_s < 1.5$ and restricting to sources with reliable shape measurements, we obtained a total number density of $\sim 17 \text{ arcmin}^{-2}$. We split source galaxies into four tomographic redshift bins with edges [0.3, 0.6, 0.9, 1.2, 1.5] and construct convergence maps via Kaiser-Squires inversion [77].

B. PDF and power spectrum

The primary summary statistic we consider is the lensing PDF, calculated as histograms of pixels in normalized convergence maps. One crucial step in our computation is that we define the PDF over the signal-to-noise ratio—i.e., each convergence map is divided by its standard deviation before its pixels are histogrammed. This is important in order to remove information duplicated in the power spectrum. The PDF’s non-Gaussian character becomes apparent in Fig. 1, where the tails deviate from a Gaussian (which is dominated by galaxy shape noise).




FIG. 1. Example of the non-Gaussian lensing convergence PDFs from the HSC Y1 data (orange) and the fiducial simulation (blue) for the tomographic redshift bin $0.6 < z_s < 0.9$ and with a smoothing scale $\theta_s = 5 \text{ arcmin}$. For most of the data points, the error bars are invisible. Unfilled data points are not used in the analysis. The PDF’s deviation from a Gaussian (black) provides the cosmological information to be exploited in this work.

We measure PDFs on maps with 0.88 arcmin side-length pixels smoothed with Gaussian filters $\propto \exp(-\theta^2/\theta_s^2)$, where we choose $\theta_s = \{5, 7, 10\} \text{ arcmin}$ to mitigate systematics while retaining non-Gaussian information. A Gaussian filter has optimal joint localization in configuration and reciprocal space and is thus a common choice [e.g., [17,47,48]]. The presence of masking (where we set shear to zero) means that the filter may decrease the signal-to-noise ratio slightly, but no biases can arise, since the mask is consistently implemented in data and simulations. For each smoothing scale, we histogram in 19 equally spaced bins between -4 and 4 (cf. Fig. 1). Of these, we remove the first three bins to minimize contamination by baryons according to tests described in Sec. II E [also, cf. [78]]. Finally, we remove the tenth bin, as otherwise, the linear constraint from fixing the first three moments would render the PDF’s auto-covariance almost singular.

We also consider the auto-power spectrum for each tomographic redshift bin, deconvolving the survey mask using the pseudo- C_ℓ method as implemented in NaMaster [79,80]. We measure C_ℓ^{xx} in 14 logarithmically spaced bins in angular multipole $81 < \ell < 6580$. Of these, we remove $\ell < 285$, as Ref. [81] found unmodeled systematic errors for these scales in HSC Y1 data. Furthermore, we remove $\ell > 1000$ due to possible contamination by baryons. We note that $\ell_{\text{max}} \sim 1000$ is also comparable to the minimum smoothing scale of 5 arcmin used in the PDF data vector. This leaves us with four ℓ bins. In contrast to previous two-point analyses [82,83] we do not include cross spectra between tomographic redshift bins. Upon unblinding the power-spectrum-only posteriors, we found that the highest tomographic redshift bin causes significant shifts in S_8 .³ Thus, we exclude the highest-redshift bin in our analysis.

³This may be due to deficiencies in our simulations; for example, the simulations use lower resolution at high redshifts, which may not be sufficient. Alternatively, the highest source redshift bin may be subject to photo- z calibration errors [7,8].

In summary, our raw data vector consists of the PDF with 135 dimensions (3 smoothing scales \times 3 redshifts \times 15 bins) and the power spectrum with 12 dimensions (3 redshifts \times 4 bins).

C. Simulations

We adopt two sets of N -body simulations in our analysis:

- (1) To model the covariance, we use 2268 realizations of the HSC Y1 footprint at a fiducial point ($S_8 = 0.791$, $\Omega_m = 0.279$), generated from the 108 quasi-independent [84] full-sky maps of Ref. [85].
- (2) To model our statistics, we use simulations at 100 different cosmological models with varying values of S_8 and Ω_m , as introduced in Ref. [86]. For each cosmology, 50 quasi-independent realizations are generated by randomly placing observers within the periodic box. These simulations are used to construct emulators for the mean theory prediction; thus, they are designed to have better accuracy than the fiducial simulations.

For each of these 7768 ($= 2268 + 100 \times 50$) realizations, mock galaxy shapes are generated following Refs. [87–89]. Convergence map generation and summary statistic measurements are performed on these simulated mocks identically to the processes used for the real data.

D. Likelihood and inference

For the PDF, to reduce the size of the data vector and Gaussianize its likelihood, we score-compress the logarithmic PDF under the approximation of a Gaussian likelihood [MOPED, [90]].⁴ This reduces the number of PDF bins from 135 to only 2, corresponding to the number of parameters (S_8 and Ω_m). We then construct at each of the 100 cosmological models a 2×2 covariance matrix, using the 50 realizations. Finally, we build emulators of both the compressed PDF and the cosmology-dependent inverse covariance using Gaussian processes.⁵ If we use a cosmology-independent covariance, the S_8 -posterior tightens, and its peak is almost unchanged. However, with that choice, the ranks plot discussed below would indicate an overconfident posterior.

The power spectrum data vector is small enough, and its distribution is known to be close to Gaussian. Therefore, we apply no data compression to the power spectrum. We build a Gaussian process emulator to model the power spectrum and estimate a cosmology-independent

⁴We note that upon combining the PDF with the power spectrum, the compression is likely suboptimal; future work could investigate more rigorous ways to preserve a maximum of information. Of course, the choice of likelihood to compute the score cannot induce biases, only information loss.

⁵Analytic methods to compute the covariance have been presented in Refs. [66,91].




FIG. 2. Validations of our likelihoods. The histograms are over hundreds of Markov chains, run with realizations from the cosmo-varied simulation set as “observation.” The x axis is the position of the simulation’s true S_8 value in the ordered chain. A uniform distribution is expected for a valid likelihood. The gray band represents the 2σ interval, approximating the bin counts as Poissonian.

covariance matrix using the fiducial simulations. To jointly analyze the compressed PDF and the power spectrum, we approximate the cross-correlation between them as constant and estimate it from the fiducial simulations.

We adopt uniform priors on our parameters S_8 and Ω_m , in the intervals [0.5, 1.0] and [0.2, 0.4], respectively. Our prior is well covered by the available simulations. Markov chain Monte Carlo (MCMC) sampling is performed using EMCEE [92,93].

We validate our likelihoods with a “ranks plot” [e.g., [94]], shown in Fig. 2. The rank statistic is a simple calibration test which relies on the fact that for a valid likelihood, random draws from the posterior should be statistically indistinguishable from the true parameter vector. To construct the ranks plot, we run MCMC on realizations drawn from the 30 cosmology-varied simulations within our prior. To cleanly separate our training and test sets, all emulators are rebuilt without the test cosmology. We then order each Markov chain by its S_8 values and find the rank of the true S_8 . For a valid likelihood, it should be impossible to statistically distinguish the true S_8 from randomly selected Monte Carlo samples, so the ranks should follow a uniform distribution. If the posterior is overconfident (underconfident), the histogram exhibits a U-shape (inverse U). However, because we only have a relatively small number of simulated cosmologies available within our prior, a perfectly uniform distribution may not be possible to attain. Indeed, Fig. 2 shows an approximately uniform distribution for all data vector choices (PDF, power spectrum, and the two of them jointly), except for spikes at the edges for the power spectrum and joint statistic. These spikes are attributable to a small number of simulations near the prior boundary, where the power spectrum emulation appears to work less well. We do not expect these to affect our results, as they have little overlap with our final posterior, and the problem is mostly due to underestimated tails, which are less important for the confidence levels typically considered.




FIG. 3. Left: estimate of systematic errors. We contaminate the PDF + $C_{\ell}^{\kappa\kappa}$ data vector with various systematics as described in the text. No systematic exceeds $0.45\sigma(S_8)$. Right: assessment of training set and emulator quality. Each point is one of the cosmologies for which cosmo-varied realizations exist. The bias indicated is estimated by averaging over chains run on 25 realizations, while the emulator is trained with the other 25 realizations, and without including the cosmology where the inference is performed.

E. Systematics

We study potential biases caused by effects not included or wrongly described in the simulated model by running inference on fiducial mocks contaminated with the following methods:

- (1) We simulate miscalibration of multiplicative bias by shifting it by $\pm 1\%$, corresponding to the uncertainty on the mean quoted in Ref. [95].
- (2) We assess the impact of photometric redshift uncertainty by generating fiducial mocks with source redshifts from two other codes, FRANKENZ and MIZUKI.
- (3) We model the impact of baryonic effects using the κ TNG simulations [96]—two sets of convergence maps with and without baryons. We contaminate the fiducial mocks with the ratio of hydrodynamic to dark-matter-only data vectors.
- (4) We estimate the impact of intrinsic alignments similarly to the baryons’ case by multiplying the fiducial data vectors with a fixed ratio which we obtain from mocks with nonlinear intrinsic alignments [97]. We choose alignment amplitudes, $A_{IA} = \{-0.32, 1.18\}$, comparable to the 1σ constraints in Refs. [82,83].

We show the resulting biases on S_8 for these systematics in the left panel of Fig. 3 (also see the more comprehensive discussion in Ref. [74]). The most constraining joint PDF + $C_{\ell}^{\kappa\kappa}$ data vector is used. We do not find biases exceeding 0.45σ in S_8 .

In addition, we also consider potential biases caused by imperfections in the simulations and emulators:

- (1) The mean recovered S_8 on the fiducial simulations is about 1σ higher than the input, indicating a systematic difference with the cosmology-varied

simulations. This is likely due to the lower fidelity of the fiducial simulations. With this in mind, the observed 1σ bias can be considered an upper bound on biases caused by resolution effects in the N -body simulations.

- (2) To assess biases caused by the limited training set size, at each cosmology we divide the realizations into 25 samples each for training and testing. We match the realization indices (i.e., observer positions) between cosmologies so as to maximize statistical independence between training and testing samples. In addition, all realizations at the test cosmology are removed from the training. The right panel of Fig. 3 shows the average biases on S_8 within our prior range, in units of the standard deviation. The bias is computed using the mode of the marginalized S_8 posterior to minimize the effect of the prior, which can shift the mean. The bias is a few tens of percent for most of the points, except at the edges, where the emulator quality decreases. Since the described test aggressively reduces the training set size, our bias estimates are likely a conservative upper bound.

F. Blinding

To help build best practices for non-Gaussian statistical analyses with Stage-IV surveys [e.g., [48]], we follow a three-step blinding procedure in our analysis:

- (1) We build the mock generation pipeline using one random realization from the fiducial model as “observation.” This includes shear bias correction, convergence map reconstruction, masking, in-painting, and summary statistics measurements.
- (2) We construct the inference pipeline using the ranks plot and similar tests. This includes data compression, covariances, emulators, and MCMC sampling.
- (3) We select smoothing scales, redshift bins, and cuts on data vectors to minimize the impact of systematics.
- (4) First unblinding: B-modes. We compare the power spectra and PDFs measured in B-mode maps of the HSC Y1 data and our fiducial simulation.⁶ We validate that the two distributions are statistically consistent.
- (5) Second unblinding: power spectrum. We apply inference to the measured HSC Y1 power spectrum and compare our results internally between different redshift bins and to the official HSC Y1 analyses by Refs. [82,83]. We discover systematics in the highest tomographic redshift bin and remove it from our final analysis.

⁶B-mode maps are built by rotating galaxy shapes by 45 degrees.

- (6) Third unblinding: PDF. We unblinded the PDF-only posterior during an invited presentation at the HSC weak lensing working group telecon. The initial posterior was approximately uniform. This failure was due to a bug in computing the standard deviations of the HSC data maps.⁷ The bug did not affect our previous steps, so once we resolved it, we did not change any other parts of the pipeline.
- (7) After the initial submission of this paper, we learned that there are two additional multiplicative bias parameters arising from galaxy size selection and redshift-dependent responsivity corrections. Post-unblinding, we updated our analysis to account for this total multiplicative bias. The correction is very small and hence did not change our conclusions. However, it made our power-spectrum-only results more consistent with the official analysis [82].

Since the HSC Y1 data are already public, our blinding procedure is rather an honor system.⁸

III. RESULTS

We show the S_8 posteriors obtained from the power spectrum, the PDF, and the two combined in Fig. 4. In all cases, the constraint on Ω_m is prior dominated, and thus is not shown here for brevity. For reference, we also show the 68% intervals from the official HSC Y1 two-point analyses with the power spectrum [82] and the two-point correlation function [83].

Our two-point result (blue) agrees well with the previous analyses in terms of mean value. However, our posterior is wider: $\sigma \simeq 0.042$, compared to $\sigma \simeq 0.03$ in the previous two-point analyses. This is expected, as these analyses differ in many respects, such as additional cross spectra between tomographic redshift bins, choices of priors, treatment of systematic errors, data vector binning, and scale cuts.

Our PDF-only posterior (orange) shows a clear detection of S_8 , or specifically *nonlinear* clustering, at 9σ . This is after removing variance information from the PDF and thus is almost independent of the two-point results. The mean S_8 value is slightly higher than the one from C_ℓ^{KK} , but it is statistically consistent. Upon combining the PDF and C_ℓ^{KK} , we observe a slight upwards shift in the mean S_8 and a

⁷Thanks to the feedback from the HSC weak lensing working group, we learned a valuable lesson of comparing the PDFs before moving on to inferences. They also pointed out that a more sophisticated blinding policy, as typically employed by collaborations, would have flagged this issue before unblinding.

⁸During our study, we also became aware of other recently completed non-Gaussian statistical analyses using the same dataset [47,98]. To avoid confirmation bias, we refrained from reading these papers until we finalized our results.




FIG. 4. S_8 posteriors from different data vector choices. Since the posteriors are slightly asymmetric, we quote their modes as well as the symmetric 68% confidence interval. For comparison, we include 68% intervals from Refs. [82,83]. Note that these analyses differ in various details. The inset compares the best-fit model to the data (PTE = 0.86). We remind the reader that the PDF does not contain information on the variance.

$\approx 10\%$ tightening of the posterior. This shows that the PDF indeed contains complementary information to the power spectrum.

In many aspects, our results match the CNN analysis from Ref. [47], including the S_8 posterior and an $\mathcal{O}(10\%)$ improvement compared to the power spectrum.

IV. CONCLUSIONS

In this work, we obtained the first cosmological constraints using the probability distribution function of weak lensing convergence. We built convergence maps using the HSC Y1 shear catalogue, constructed a likelihood based on tailored numerical simulations, and validated that known survey and astrophysical systematics are under control after applying cuts on smoothing scales, redshift, and data vectors. We designed and followed a three-step blinding procedure to minimize confirmation bias. We make our inference code publicly available at this URL [99].

The PDF improves the power-spectrum-only constraint by about 10%. For the clustering amplitude, we obtain $S_8 = 0.860_{-0.109}^{+0.066}$ and $S_8 = 0.798_{-0.042}^{+0.029}$ from the lensing PDF alone and the combination of PDF and power spectrum, respectively (68% CL). We computed the PDF on convergence maps normalized by their standard deviation; hence, we maximally removed two-point information from the PDF. Our results are consistent with previous analyses on the same data and show that the PDF provides additional information not contained in the power spectrum.

We find no tension between the S_8 inferred from HSC Y1 lensing and from primordial CMB measurements.

Future work could investigate alternatives to the probably lossy compression step, as well as the forward modeling of systematics to allow inclusion of smaller scales. Additionally, Stage-IV data will be sensitive to cosmological parameters beyond S_8 —for example, the neutrino mass sum. The PDF could be instrumental in complementing two-point statistics in order to optimally constrain such model extensions.

ACKNOWLEDGMENTS

We thank Sihao Cheng, Will Coulton, Daniela Grandón, Kevin Huffenberger, Xiangchong Li, Surhud More, Ken Osato, David Spergel, Sunao Sugiyama, and Masahiro Takada for useful discussions. We thank Joachim Harnois-Déraps for sharing the IA mocks. We thank the organizers of the Kyoto CMBxLSS workshop, during which this work was completed. The work of L. T. is supported by NSF Grant No. AST 2108078. This work was supported by JSPS KAKENHI Grants No. 23K13095 and No. 23H00107 (to J. L.), and by Grants No. 19K14767 and No. 20H05861 (to M. S.). This work was supported in part by the UTokyo-Princeton Strategic Partnership Teaching and Research Collaboration. This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. The authors are pleased to acknowledge that the work reported on in this paper was substantially performed using the Princeton Research Computing resources at Princeton University, which is a consortium of groups led by the Princeton Institute for Computational Science and Engineering (PICSciE) and the Office of Information Technology’s Research Computing.

APPENDIX A: INFORMATION

In Fig. 5, we illustrate where the information in the PDF is primarily coming from.

In the left panel, we focus on redshift bins and smoothing scales. We observe that, analogously to the two-point statistics, the highest redshift bin is the most informative. Furthermore, the largest smoothing scale (10 arcmin) is quite uninformative, due to the nearly Gaussian character of the field. On the other hand, the smaller smoothing scales contain information, and the compression appears to use the difference between them.

In the right panel, we focus on the convergence bins. Even though it might be expected that the higher signal-to-noise ratio close to the PDF’s peak would make these bins the most informative, we observe that the high- κ tail is actually where most of the information is coming from. This is consistent with our picture in Fig. 1, where the deviations from Gaussianity in the tail are the focus of this work.

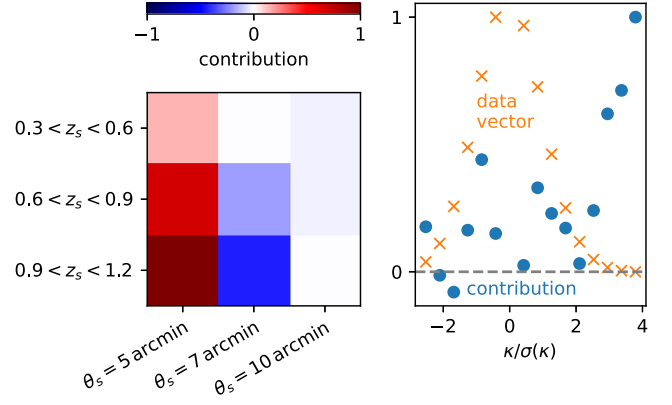


FIG. 5. Illustration of the compression weights for the PDF. We picked a high- S_8 cosmology (with close to fiducial Ω_m) from our simulations. By taking the elementwise product between the first MOPED compression vector (corresponding to S_8) and the data vector, we can gain intuition as to where information is coming from. For this figure, we subtract the corresponding product for the fiducial simulations. Note that the data vector shown in orange is just for illustration and has been rescaled and shifted for convenience. Left: contributions from different source redshifts and smoothing scales (with the κ -bins summed over). Right: contributions from different κ -bins with source redshifts and smoothing scales summed over.

APPENDIX B: DATA VECTOR CHOICES

In Fig. 6, we explore the effects of various modifications to the data vector. First, we remove the lowest and highest tomographic redshift bins (orange, green) and observe that the higher one contributes much more of the signal-to-noise ratio. Then, we add the tomographic redshift bin that was excluded in our baseline analysis, cf. Sec. II B (red). As discussed in the text, including this bin leads to a substantial shift in the posterior.

The last three tests apply only to the PDF part of the data vector. First, we remove the smallest smoothing scale




FIG. 6. The effect of various cuts and extensions in the data vector on the final posterior. All posteriors plotted are for PDF + $C_{\ell}^{\kappa\kappa}$, the blue baseline being the same curve as in Fig. 4.

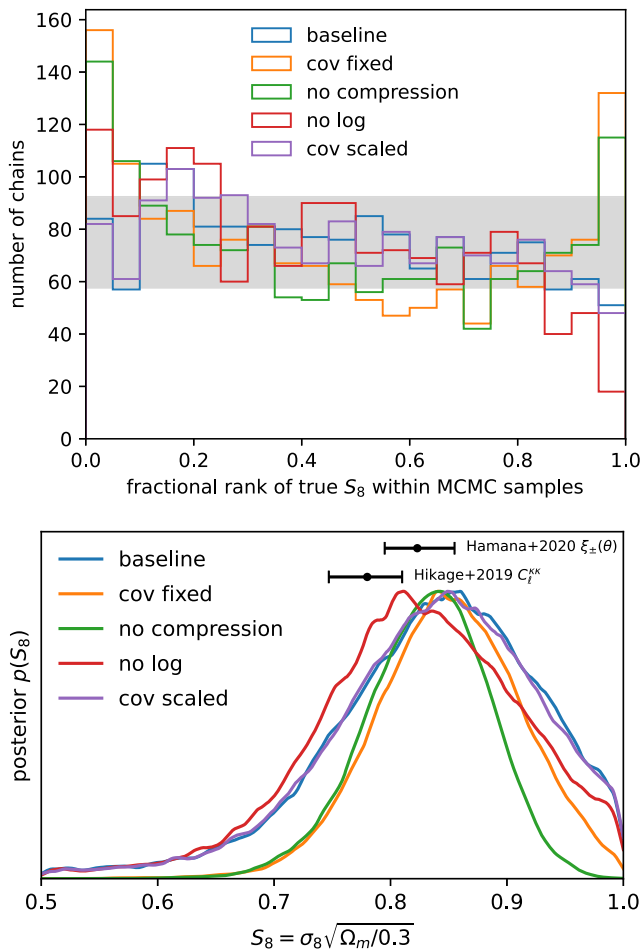


FIG. 7. The effect of changing the PDF likelihood in various ways (cf. Sec. II D). In all plots, the data vector considered is PDF only. The top panel is the ranks plot, cf. Fig. 2, and the bottom panel shows the posteriors on real data obtained with the different likelihoods. We observe that the baseline likelihood used in the main part of the paper exhibits the healthiest calibration diagnostics. Furthermore, we see that the different likelihood assumptions lead to visible but mild shifts in the posterior.

(purple). Then, we add a 2 arcmin smoothing scale (brown). We observe that using the smaller smoothing scale tightens the posterior substantially, indicating that there is useful information. However, in our baseline analysis, such small scales are not used due to concerns about baryonic systematics. Finally, we delete the 17th convergence bin (instead of the tenth, cf. Sec. II B), and we find that this modification of the data vector does not alter the posterior.




FIG. 8. Impact of tomography in the PDF-only analysis. We consider only the lower three redshift bins, as explained in Sec. II B. We observe that the posterior obtained with tomography (blue) is tighter than the one obtained without it (orange). The combined posterior (green) is consistent with the individual ones but slightly broader than the tomographic result, presumably due to small inefficiencies in the compression.

APPENDIX C: LIKELIHOOD CHOICES

In Fig. 7, we show what happens to the PDF-only posterior when various choices in our likelihood construction are altered. As shown in the upper panel, the baseline used in the main text (blue) exhibits the best rank statistics (closest to uniform distribution). The modifications deviate from uniform distribution. The resulting posteriors are still broadly consistent.

APPENDIX D: TOMOGRAPHY

For two-point analyses, including the auto- and cross-power spectra from tomographic redshift bins saturates the information. The situation is more complicated with non-Gaussian statistics, because the shape noise propagates in a nonlinear way into the data vector and posteriors. In the main text, we used a tomographic analysis, motivated by previous work that indicated benefits to tomography [65]. In Fig. 8, we show what happens when we collapse all sources in a single bin (orange). We observe that the tomographic analysis does yield a somewhat tighter posterior. Of course, one can also combine tomographic and nontomographic data vectors, or use various other mergers of redshift bins. We do not expect such strategies to significantly affect the constraining power in our case, but more generally, they warrant further study.

- [1] G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett *et al.*, *Astrophys. J. Suppl. Ser.* **208**, 19 (2013).
- [2] Planck Collaboration, *Astron. Astrophys.* **641**, A6 (2020).
- [3] S. Aiola, E. Calabrese, L. Maurin, S. Naess, B. L. Schmitt *et al.*, *J. Cosmol. Astropart. Phys.* **12** (2020) 047.
- [4] M. Asgari, C.-A. Lin, B. Joachimi, B. Giblin, C. Heymans *et al.*, *Astron. Astrophys.* **645**, A104 (2021).
- [5] A. Amon, D. Gruen, M. A. Troxel, N. MacCrann, S. Dodelson *et al.*, *Phys. Rev. D* **105**, 023514 (2022).
- [6] L. F. Secco, S. Samuroff, E. Krause, B. Jain, J. Blazek, M. Raveri *et al.*, *Phys. Rev. D* **105**, 023515 (2022).
- [7] R. Dalal, X. Li, A. Nicola, J. Zuntz, M. A. Strauss *et al.*, arXiv:2304.00701.
- [8] X. Li, T. Zhang, S. Sugiyama, R. Dalal, M. M. Rau *et al.*, arXiv:2304.00702.
- [9] B. Jain and L. Van Waerbeke, *Astrophys. J. Lett.* **530**, L1 (2000).
- [10] L. Marian, R. E. Smith, and G. M. Bernstein, *Astrophys. J. Lett.* **698**, L33 (2009).
- [11] J. P. Dietrich and J. Hartlap, *Mon. Not. R. Astron. Soc.* **402**, 1049 (2010).
- [12] J. M. Kratochvil, Z. Haiman, and M. May, *Phys. Rev. D* **81**, 043519 (2010).
- [13] J. Liu, A. Petri, Z. Haiman, L. Hui, J. M. Kratochvil, and M. May, *Phys. Rev. D* **91**, 063507 (2015).
- [14] A. Peel, C.-A. Lin, F. Lanusse, A. Leonard, J.-L. Starck, and M. Kilbinger, *Astron. Astrophys.* **599**, A79 (2017).
- [15] N. Martinet, P. Schneider, H. Hildebrandt, H. Shan, M. Asgari *et al.*, *Mon. Not. R. Astron. Soc.* **474**, 712 (2018).
- [16] J. Harnois-Déraps, N. Martinet, T. Castro, K. Dolag, B. Giblin, C. Heymans, H. Hildebrandt, and Q. Xia, *Mon. Not. R. Astron. Soc.* **506**, 1623 (2021).
- [17] A. Sabyr, Z. Haiman, J. M. Z. Matilla, and T. Lu, *Phys. Rev. D* **105**, 023505 (2022).
- [18] C. T. Davies, M. Cautun, B. Giblin, B. Li, J. Harnois-Déraps, and Y.-C. Cai, *Mon. Not. R. Astron. Soc.* **507**, 2267 (2021).
- [19] D. Munshi, L. van Waerbeke, J. Smidt, and P. Coles, *Mon. Not. R. Astron. Soc.* **419**, 536 (2012).
- [20] J. M. Kratochvil, E. A. Lim, S. Wang, Z. Haiman, M. May, and K. Huffenberger, *Phys. Rev. D* **85**, 103513 (2012).
- [21] A. Petri, Z. Haiman, L. Hui, M. May, and J. M. Kratochvil, *Phys. Rev. D* **88**, 123002 (2013).
- [22] A. Petri, J. Liu, Z. Haiman, M. May, L. Hui, and J. M. Kratochvil, *Phys. Rev. D* **91**, 103511 (2015).
- [23] M. Vicinanza, V. F. Cardone, R. Maoli, R. Scaramella, X. Er, and I. Tereno, *Phys. Rev. D* **99**, 043534 (2019).
- [24] G. A. Marques, J. Liu, J. M. Z. Matilla, Z. Haiman, A. Bernui, and C. P. Novaes, *J. Cosmol. Astropart. Phys.* **06** (2019) 019.
- [25] C. Parroni, V. F. Cardone, R. Maoli, and R. Scaramella, *Astron. Astrophys.* **633**, A71 (2020).
- [26] J. Feldbrugge, M. van Engelen, R. van de Weygaert, P. Pranav, and G. Vegter, *J. Cosmol. Astropart. Phys.* **09** (2019) 052.
- [27] C. Parroni, É. Tollet, V. F. Cardone, R. Maoli, and R. Scaramella, *Astron. Astrophys.* **645**, A123 (2021).
- [28] S. Heydenreich, B. Brück, and J. Harnois-Déraps, *Astron. Astrophys.* **648**, A74 (2021).
- [29] S. Heydenreich, B. Brück, P. Burger, J. Harnois-Déraps, S. Unruh, T. Castro, K. Dolag, and N. Martinet, *Astron. Astrophys.* **667**, A125 (2022).
- [30] S. Pires, J. L. Starck, A. Amara, A. Réfrégier, and R. Teyssier, *Astron. Astrophys.* **505**, 969 (2009).
- [31] S. Cheng, Y.-S. Ting, B. Ménard, and J. Bruna, *Mon. Not. R. Astron. Soc.* **499**, 5902 (2020).
- [32] S. Cheng and B. Ménard, arXiv:2112.01288.
- [33] M. Takada and B. Jain, *Mon. Not. R. Astron. Soc.* **337**, 875 (2002).
- [34] M. Kilbinger and P. Schneider, *Astron. Astrophys.* **442**, 69 (2005).
- [35] L. Van Waerbeke, J. Benjamin, T. Erben, C. Heymans, H. Hildebrandt *et al.*, *Mon. Not. R. Astron. Soc.* **433**, 3373 (2013).
- [36] L. Porth and R. E. Smith, *Mon. Not. R. Astron. Soc.* **508**, 3474 (2021).
- [37] M. Gatti, B. Jain, C. Chang, M. Raveri, D. Zürcher, L. Secco, L. Whiteway, N. Jeffrey, C. Doux, T. Kacprzak, D. Bacon, P. Fosalba *et al.*, *Phys. Rev. D* **106**, 083509 (2022).
- [38] A. Cooray and W. Hu, *Astrophys. J.* **548**, 7 (2001).
- [39] M. Takada and B. Jain, *Mon. Not. R. Astron. Soc.* **344**, 857 (2003).
- [40] M. Takada and B. Jain, *Mon. Not. R. Astron. Soc.* **348**, 897 (2004).
- [41] S. Dodelson and P. Zhang, *Phys. Rev. D* **72**, 083001 (2005).
- [42] D. Huterer, M. Takada, G. Bernstein, and B. Jain, *Mon. Not. R. Astron. Soc.* **366**, 101 (2006).
- [43] S. Vafaei, T. Lu, L. van Waerbeke, E. Semboloni, C. Heymans, and U.-L. Pen, *Astropart. Phys.* **32**, 340 (2010).
- [44] D. Gruen, O. Friedrich, E. Krause, J. DeRose, R. Cawthon, C. Davis, J. Elvin-Poole, E. S. Rykoff, R. H. Wechsler *et al.*, *Phys. Rev. D* **98**, 023507 (2018).
- [45] J. Fluri, T. Kacprzak, A. Lucchi, A. Refregier, A. Amara, T. Hofmann, and A. Schneider, *Phys. Rev. D* **100**, 063514 (2019).
- [46] J. Fluri, T. Kacprzak, A. Lucchi, A. Schneider, A. Refregier, and T. Hofmann, *Phys. Rev. D* **105**, 083518 (2022).
- [47] T. Lu, Z. Haiman, and X. Li, *Mon. Not. R. Astron. Soc.* **521**, 2050 (2023).
- [48] Euclid Collaboration, *Astron. Astrophys.* **675**, A120 (2023).
- [49] N. Porqueres, A. Heavens, D. Mortlock, and G. Lavaux, *Mon. Not. R. Astron. Soc.* **509**, 3194 (2022).
- [50] N. Porqueres, A. Heavens, D. Mortlock, G. Lavaux, and T. L. Makinen, arXiv:2304.04785.
- [51] A. Bayer, U. Seljak, and C. Modi, in *Machine Learning for Astrophysics, Workshop at the Fortieth International Conference on Machine Learning (ICML 2023)* (2023), p. 3.
- [52] B. Dai and U. Seljak, in *Machine Learning for Astrophysics, Workshop at the Fortieth International Conference on Machine Learning (ICML 2023)* (2023), p. 10.
- [53] D. Zürcher, J. Fluri, V. Ajani, S. Fischbacher, A. Refregier, and T. Kacprzak, *Mon. Not. R. Astron. Soc.* **525**, 761 (2023).
- [54] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta *et al.*, *Astrophys. J.* **873**, 111 (2019).
- [55] R. Laureijs, J. Amiaux, S. Arduini, J. L. Auguères, J. Brinchmann, R. Cole *et al.*, arXiv:1110.3193.

- [56] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge *et al.*, [arXiv:1503.03757](https://arxiv.org/abs/1503.03757).
- [57] P. Valageas, *Astron. Astrophys.* **354**, 767 (2000).
- [58] P. Valageas, *Astron. Astrophys.* **356**, 771 (2000).
- [59] D. Munshi and B. Jain, *Mon. Not. R. Astron. Soc.* **318**, 109 (2000).
- [60] F. Bernardeau and P. Valageas, *Astron. Astrophys.* **364**, 1 (2000).
- [61] H. Aihara, N. Arimoto, R. Armstrong, S. Arnouts, N. A. Bahcall *et al.*, *Publ. Astron. Soc. Jpn.* **70**, S4 (2018).
- [62] H. Aihara, R. Armstrong, S. Bickerton, J. Bosch, J. Coupon *et al.*, *Publ. Astron. Soc. Jpn.* **70**, S8 (2018).
- [63] J. Liu, J. C. Hill, B. D. Sherwin, A. Petri, V. Böhm, and Z. Haiman, *Phys. Rev. D* **94**, 103501 (2016).
- [64] K. Patton, J. Blazek, K. Honscheid, E. Huff, P. Melchior, A. J. Ross, and E. Suchyta, *Mon. Not. R. Astron. Soc.* **472**, 439 (2017).
- [65] J. Liu and M. S. Madhavacheril, *Phys. Rev. D* **99**, 083508 (2019).
- [66] L. Thiele, J. C. Hill, and K. M. Smith, *Phys. Rev. D* **102**, 123545 (2020).
- [67] A. Boyle, C. Uhlemann, O. Friedrich, A. Barthelemy, S. Codis, F. Bernardeau, C. Giocoli, and M. Baldi, *Mon. Not. R. Astron. Soc.* **505**, 2886 (2021).
- [68] B. Giblin, Y.-C. Cai, and J. Harnois-Déraps, *Mon. Not. R. Astron. Soc.* **520**, 1721 (2023).
- [69] P. Reimberg and F. Bernardeau, *Phys. Rev. D* **97**, 023524 (2018).
- [70] A. Barthelemy, S. Codis, C. Uhlemann, F. Bernardeau, and R. Gavazzi, *Mon. Not. R. Astron. Soc.* **492**, 3420 (2020).
- [71] A. Barthelemy, S. Codis, and F. Bernardeau, *Mon. Not. R. Astron. Soc.* **494**, 3368 (2020).
- [72] K. Kainulainen and V. Marra, *Phys. Rev. D* **84**, 063004 (2011).
- [73] K. Kainulainen and V. Marra, *Phys. Rev. D* **83**, 023009 (2011).
- [74] G. A. Marques, J. Liu, M. Shirasaki, L. Thiele, D. Grandón, K. M. Huffenberger, S. Cheng, J. Harnois-Déraps, K. Osato, and W. R. Coulton, [arXiv:2308.10866](https://arxiv.org/abs/2308.10866).
- [75] R. Mandelbaum, H. Miyatake, T. Hamana, M. Oguri, M. Simet *et al.*, *Publ. Astron. Soc. Jpn.* **70**, S25 (2018).
- [76] M. Tanaka, J. Coupon, B.-C. Hsieh, S. Mineo, A. J. Nishizawa *et al.*, *Publ. Astron. Soc. Jpn.* **70**, S9 (2018).
- [77] N. Kaiser and G. Squires, *Astrophys. J.* **404**, 441 (1993).
- [78] J. Sunseri, Z. Li, and J. Liu, *Phys. Rev. D* **107**, 023514 (2023).
- [79] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, *Astrophys. J.* **567**, 2 (2002).
- [80] D. Alonso Monge, J. Sanchez, and A. Slosar, *Mon. Not. R. Astron. Soc.* **484**, 4127 (2019).
- [81] M. Oguri, S. Miyazaki, C. Hikage, R. Mandelbaum, Y. Utsumi *et al.*, *Publ. Astron. Soc. Jpn.* **70**, S26 (2018).
- [82] C. Hikage, M. Oguri, T. Hamana, S. More, R. Mandelbaum *et al.*, *Publ. Astron. Soc. Jpn.* **71**, 43 (2019).
- [83] T. Hamana, M. Shirasaki, S. Miyazaki, C. Hikage, M. Oguri *et al.*, *Publ. Astron. Soc. Jpn.* **72**, 16 (2020).
- [84] A. Petri, Z. Haiman, and M. May, *Phys. Rev. D* **93**, 063524 (2016).
- [85] R. Takahashi, T. Hamana, M. Shirasaki, T. Namikawa, T. Nishimichi, K. Osato, and K. Shiroyama, *Astrophys. J.* **850**, 24 (2017).
- [86] M. Shirasaki, K. Moriwaki, T. Oogi, N. Yoshida, S. Ikeda, and T. Nishimichi, *Mon. Not. R. Astron. Soc.* **504**, 1825 (2021).
- [87] M. Shirasaki and N. Yoshida, *Astrophys. J.* **786**, 43 (2014).
- [88] M. Shirasaki, M. Takada, H. Miyatake, R. Takahashi, T. Hamana, T. Nishimichi, and R. Murata, *Mon. Not. R. Astron. Soc.* **470**, 3476 (2017).
- [89] M. Shirasaki, T. Hamana, M. Takada, R. Takahashi, and H. Miyatake, *Mon. Not. R. Astron. Soc.* **486**, 52 (2019).
- [90] A. F. Heavens, R. Jimenez, and O. Lahav, *Mon. Not. R. Astron. Soc.* **317**, 965 (2000).
- [91] C. Uhlemann, O. Friedrich, A. Boyle, A. Gough, A. Barthelemy, F. Bernardeau, and S. Codis, *Open J. Astrophys.* **6**, 1 (2023).
- [92] D. Foreman-Mackey, A. Conley, W. Meierjurgen Farr, D. W. Hogg, D. Lang, P. Marshall, A. Price-Whelan, J. Sanders, and J. Zuntz, EMCEE: The MCMC Hammer, Astrophysics Source Code Library, record ascl:1303.002 (2013), ascl:1303.002.
- [93] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [94] A. Vehtari, A. Gelman, D. Simpson, B. Carpenter, and P.-C. Bürkner, *Bayesian Anal.* **16**, 667 (2021).
- [95] R. Mandelbaum, F. Lanusse, A. Leauthaud, R. Armstrong, M. Simet, H. Miyatake, J. E. Meyers, J. Bosch, R. Murata, S. Miyazaki, and M. Tanaka, *Mon. Not. R. Astron. Soc.* **481**, 3170 (2018).
- [96] K. Osato, J. Liu, and Z. Haiman, *Mon. Not. R. Astron. Soc.* **502**, 5593 (2021).
- [97] J. Harnois-Déraps, N. Martinet, and R. Reischke, *Mon. Not. R. Astron. Soc.* **509**, 3868 (2022).
- [98] X. Liu, S. Yuan, C. Pan, T. Zhang, Q. Wang, and Z. Fan, *Mon. Not. R. Astron. Soc.* **519**, 594 (2023).
- [99] https://github.com/leanderthiele/HSC_Y1_kappapdf_production.